MASSACHUSETTS INSTITUTE OF TECHNOLOGY

ARTIFICIAL INTELLIGENCE LABORATORY

A. I. Memo 359

June, 1976

SPATIAL KNOWLEDGE

Benjamin Kuipers

ABSTRACT

This paper introduces a model of spatial cognition to describe the states of partial knowledge that people have about the spatial structure of a large-scale environment. Spatial knowledge has several different representations, each of which captures one aspect of the geography. With knowledge stored in multiple representations, we must examine the procedures for assimilating new information, for solving problems, and for communicating information between representations. The model centers on an abstract machine called the TOUR machine, which executes a description of the route to drive the "You Are Here" pointer (a small working memory) through a map that describes the geography. Representations for local and global spatial knowledge are discussed in detail. The model is compared with a survey of the psychological literature. Finally, the directions of necessary and desirable future research are outlined.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Project Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643.

"Puluwat is one of a long chain of islands covering over a thousand miles of the Pacific Ocean north of New Guinea. The sailing canoe is at the very heart of the Puluwat way of life, and skilled navigators occupy the positions of highest status. There is quite a bit of drinking among men of Puluwat. A trip to Pikelot island, one mile wide and over 100 miles distant, is often launched at a drinking party. Someone jumps up and says, 'I'm going to Pikelot. Who's coming with me?' The navigator determines his sailing plan only after he is at sea; at no time is an overall plan developed for the voyage. Yet, without fail, at an average speed of 4.5 knots, the 26-foot sailing canoe with a now-sober crew arrives two days later at a one-mile wide island in the open ocean (Gladwin, 1970). What is the nature of their knowledge about the sea and the sky and the islands and the canoe that makes this possible?

Most mornings, every reader of this paper will find his way to the office, less exuberantly than the voyagers of Puluwat but no less decisively. He or she will walk, ride a bicycle, take the train, or drive. The route may be long or short, simple or complex; it may be interrupted or accomplished around diversions. What system of habits, schemes, imageries, long- or short-term memories allows him to accomplish this?

Most mornings, thousands of 6-year-old children will find their way to school. There will be protections along the way---street-crossing policemen---but there will be little fear on anyone's part that the children will get lost. The children will find their way to the school building, within the building to their classroom, and after school, perhaps, to a friend's house. However, most 3-year-olds do not typically engage in the same process. Is it that they are incapable of following routes and will thus get lost? It is quite possible that these very young children can recognize landmarks in a large-scale terrain, but have limited 'route-knowledge,' and that adults are aware of this." [Siegel and White, 1975, p. 10.]

This paper introduces a model of spatial cognition to describe the states of partial knowledge that people have about the spatial structure of a large-scale environment. Spatial knowledge has several different representations, each of which captures one aspect of the geography. With knowledge stored in multiple representations, we must examine the procedures for assimilating new information, for solving problems, and for communicating information between representations.

A large-scale environment is defined as one whose structure cannot be perceived from one location. Thus, the kind of information that can be acquired is very local and fragmented,

compared to the size of the environment as a whole. The structure of this body of knowledge must be acquired through a temporal integration of observed information. When new information is acquired, typically as a partial description of a route, it must be placed in correspondence with the description of what is already known. With this correspondence, information in each kind of description can fill gaps in the other. The body of information that is known about the geography is sometimes called a "mental map." The problems which this mental map enables people to solve include finding new routes from one place to another, and orienting themselves with respect to the positions of remote (and perhaps invisible) places. The greatest skill which people have in the spatial domain is the assimilation of new information into the appropriate representation, rather than deep or complicated problem solving.

These are activities that people perform, with varying degrees of success, many times daily. This research proposes a computational model of the representations and processes that produce these kinds of human behavior. This involves distinguishing among the kinds of knowledge involved, proposing representations for those different kinds of knowledge, and showing what interactions take place between them. The first subgoal is to identify and describe a class of human behavior to which a model can be addressed. The second is to propose a model---a computational process---which can explain that behavior. These subgoals depend on each other because the vocabulary of the model is needed to identify and describe the phenomena with precision. I will treat these subgoals together through most of this paper.

The development of this model (which I call the TOUR model) has relied on three kinds of evidence: 1) a review of the psychological literature, 2) a series of qualitative experiments,

collecting protocols and drawn maps from individual subjects, and 3) the creation of a computer program to simulate the behaviors being described. I compare the TOUR model in some detail with the theory that Siegel and White (1975) derive from their survey of the psychological literature on large-scale spatial cognition, and show that the two are largely in agreement. Where they disagree, typically the TOUR model makes a precise claim, based on computational constraints, where Siegel and White remain uncommitted.

Spatial cognition presents a very rich and complex set of phenomena. This paper focusses on the representations for knowledge of large-scale space, and the processes by which new knowledge is assimilated into those representations. I am therefore omitting a number of other important issues, though I believe that the representations presented here will provide a useful basis for addressing them. One such issue is a linguistic one. It seems clear that verbal route-descriptions are generated with the intention that the hearer will be able to derive the complete intended route from the given information and his own knowledge. Only a clear description of the nature of that knowledge and the route-finding mechanism will allow us to study this process in greater depth. Another issue is that of testing potentially false information against what is already known. In the TOUR model presented below, new information is presented as English sentences, but is intended to correspond to personal observation, and thus is assumed to be correct, though local. Of course, global conclusions drawn from this local information may be misleading, distorted, or false. It seems that understanding the representation of the knowledge, and the processes by which it is acquired, is a prerequisite to understanding how new information is tested for plausibility. I am also not addressing knowledge of smaller-scale spaces: visual space and kinesthetic space.

The TOUR model predicts that only certain kinds of partial knowledge will be represented in

the map or the route. This prediction can be made precise and subjected to empirical test. Much of the detailed behavior of the model, however, is highly dependent on certain qualitative parameters in the model, and on exactly what knowledge the subject has. Thus, testing predictions about such detailed behavior poses difficult methodological problems of custom-fitting the model to each individual subject. The parameters of variation and the model of the individual are discussed again in the concluding sections of this paper.

Some readers will be concerned that the TOUR model is expressed in terms that are most appropriate to city streets, while there is a much larger variety of large-scale environments of which people have knowledge. The TOUR model includes a number of ways of describing different features of the environment. This research, like most of the literature on large-scale spaces, concentrates on the particular context of the city. The relative usefulness of the different descriptive methods will of course vary with the kind of environment, but the same methods apply to other contexts. For example, see Gladwin (1970) for an analysis of the knowledge held by navigators in the South Pacific. The navigation knowledge is very highly structured, and new knowledge is not often acquired by an experienced navigator, but the representations seem to be the same.

Other readers may be concerned that the problem of cognitive representations of spatial knowledge, or of route-finding, can be solved trivially by applying one of several simple and well-known mathematical techniques: Cartesian coordinates for position, or route-finding with an algorithm for solving the "travelling salesman" problem. (The "travelling salesman" problem is the problem of finding the shortest path through a given set of nodes in a graph.) Neither of these can be an adequate model of human spatial cognition if applied in a straight-forward way. A "Cartesian coordinates" representation for position provides a single, unrealistic

route-finding technique, but more importantly imposes an assumption of a global frame of reference for all positions. Both of these are contradicted by the simplest qualitative experiments. Similarly, "travelling salesman" algorithms presuppose a network representation which is of no help in explaining the concept of a "sense of direction", and has a computational complexity well beyond the route-finding techniques people actually use. A hybrid representation which incorporates aspects of both is responsible for explaining the interactions between them, and how knowledge is acquired and assimilated into those representations. The TOUR model is, in part, such a hybrid, and much of the research deals with the nature of the interactions between pieces of knowledge in different representations.

The model of spatial cognition that I describe involves the interaction of several kinds of knowledge, mediated by a simple machine called the TOUR machine. (Note 1.) The different kinds of knowledge are: 1) routes, represented as sequences of instructions to TOUR; 2) a small working memory, called the "You Are Here" pointer; 3) the map, which is a description of the geography; 4) the problem-solving process, which proposes routes by examining the map; and 5) a process that generates global descriptions of the map. TOUR has several novel properties which allow it to facilitate both problem-solving and the assimilation of new information.

In the next section, I will describe in some detail the TOUR machine, the "You Are Here" pointer, and the route and map representations. A detailed example of the assimilation of information from a route description into the map is included. Following that, I present several other aspects of the map, including orientation, abstraction, and more global ways of describing the geography. The next section reviews some of the relevant psychological research, and shows how the TOUR model of spatial cognition agrees with the empirical data,

and allows a more precise description of spatial knowledge than had previously been possible. A final section discusses the work yet to be done to complete the TOUR model.

THE TOUR MACHINE AND ITS DOMAIN

In this section, I describe the TOUR machine, the route programs that drive it, and the representation of the map it moves on. As I explain it, I will begin with a somewhat simplified version, which I will progressively extend to accomodate more complex phenomena.

Let us begin with a contrasting analogy. Consider a robot vehicle which can travel on a featureless plane in response to two commands:

```
FORWARD <number of units>
RIGHT <number of degrees>
```

The vehicle is completely described by a "state" that consists of its position and heading on the plane. LOGO {Note 2} is a programming system that consists of a plane, a vehicle, and an interpreter for this simple language. The key difference between LOGO and TOUR is that the featureless plane is replaced by a highly structured, network-like map. So the "state" of the traveller must be correspondingly more complex. TOUR executes a simple LOGO-like programming language that drives the "You Are Here" pointer through the map.

A PLACE is a description of a zero-dimensional geographical object: a location or landmark. A PATH is a description of a one-dimensional geographical object: often a street or a road. The PLACE description includes the PATHs that the PLACE is on, and it has a local geometry to describe the relation among those PATHs. A PATH description includes an order on the PLACEs on that PATH. This order lets us define a DIRECTION on a PATH: +1 or -1 meaning "with" or "against" the order on the PATH. The local geometry of a PLACE is an arbitrary set of radial coordinates, with respect to which the headings of (PATH DIRECTION) pairs can be defined as they leave the PLACE. These descriptions are intended to represent partial

knowledge of the geography. Therefore, the order of PLACEs on a PATH may be a partial order, and the description a PLACE has of its local geometry may not be complete. This allows the map to represent certain quite partial states of knowledge, and to assimilate new fragments of information into existing descriptions. (Note 3 summarizes the map definitions.)

The local geometry makes it possible to predict the effects of a particular turn on the state of the TOUR machine. The state of the TOUR machine is represented by the "You Are Here" pointer, which consists of: the PLACE it is at, the PATH it is on, and the DIRECTION it faces on that PATH. Just as the map can represent partial knowledge, so the "You Are Here" pointer can leave parts of the current state unspecified, and, as we shall see below, TOUR can run on an incomplete program.

A PLACE contains only part of the knowledge a human would have about a geographical place. In particular, a human has a rich sensory image, visual and otherwise, which is associated with a place and helps him recognize it. In the TOUR model, a PLACE has only that part of the knowledge which is strictly spatial in nature. Similarly, a PATH lacks information that lies apart from its spatial properties.

The basic TOUR instructions are GO-TO and TURN, which bear a strong resemblance to their LOGO counterparts, but with some important differences. First, a LOGO instruction specifies a change to be made to the assumed current state of the turtle. A fully specified TOUR instruction specifies the initial state, the change to be made, and the resulting state of the "You Are Here" pointer. Second, a TOUR instruction need not be fully specified for TOUR to be able to execute it, while an underspecified LOGO instruction has no meaning. Third, a LOGO instruction specifies an action to be executed, while the TOUR instruction specifies a state to

be achieved. The full forms of the two basic TOUR instructions are:

```
(GO-TO <from-place> <to-place> <on-path> <in-direction>)
(TURN <at-place> <from-path> <from-direction>
      <turn-amount> <to-path> <to-direction>)
```

GO-TO instructs TOUR to move the "You Are Here" pointer from one PLACE to another on the same PATH, and specifies the direction of travel with respect to the PATH order. TURN instructs TOUR to move the "You Are Here" pointer from one (PATH DIRECTION) pair to another at the same PLACE, and specifies the amount of the turn. A route is simply a sequence of instructions that TOUR can treat as a program. (Note 4 summarizes the TOUR instructions.)

The TOUR machine communicates information between several different representations of spatial knowledge by filling in underspecified instructions and descriptions when possible. When executing a fully specified instruction, TOUR matches the current position in the "You Are Here" pointer against the current state presupposed by the instruction. Then it checks the instruction for consistency with the map, and changes the "You Are Here" pointer to represent the new position. A fully specified instruction can contribute new information to the "You Are Here" pointer, or to the map, adding to the order information in a PATH, or the local geometry of a PLACE.

On the other hand, an individual TOUR instruction can be substantially underspecified. For example, the English command "Turn right" translates into a TOUR instruction lacking almost all information. The current state can be obtained from the "You Are Here" pointer, and the destination PATH and DIRECTION can be obtained by examining the local geometry of the current PLACE. Naturally, it is not always possible to fill in all the gaps in the current instruction, the "You Are Here" pointer, or the current parts of the map. TOUR can move the

"You Are Here" pointer along a route which is quite underspecified even when it is unable to fill in all the gaps, but there will be information about the geography implicit in the route description, which cannot be transferred to the map.

A route program is underspecified when the current position presupposed by an instruction is different from the "You Are Here" pointer. For example, if I am in Kendall Square, desiring a route to Harvard Square, and the proposed route is "Take Mass Ave from Central Square," the assumption is that I can find the route to Central Square. In this case, TOUR formulates a difference description, and calls the problem-solving component to propose a solution. The proposed solution to such a problem may itself be an underspecified route, with smaller gaps to be filled in their turn. A route is stored in the map, indexed under source and destination, so that it can be found when needed later.

Since the program executed by TOUR can be more fully specified than the one it was given, it produces the updated program as output, as well as having side-effects on the "You Are Here" pointer and the map. This is important because often not all of the information in the route description is assimilated into the map on a single pass. Thus, by saving the updated route description, the knowledge is available for that particular route, and later tours may be able to extract more information. They may also update the route description further.

We can now introduce two new TOUR instructions: TAKE and GET-TO. The TAKE instruction refers to a route, and lets it be used as a sub-program for a longer route. This can give a route program a hierarchical structure, and it allows frequently-used routes to be shared. The GET-TO instruction formulates a problem to be solved by the problem-solving component by specifying the state of the "You Are Here" pointer at source and destination. In case no

solution can be found, the problem-statement is left as an instruction in the route program, and TOUR continues from its destination. The forms of these instructions are:

```
(TAKE <route> <from-place> <to-place>)
(GET-TO <from-place> <from-path> <from-direction>
        <to-place> <to-path> <to-direction>).
```

The most straightforward increase in the amount of knowledge contained in the map is the addition of a new PLACE or PATH. This happens when a route program refers to the new PLACE or PATH. The new part of the map is automatically created and added by the processes that create the route. This process is typically the natural language interface which translates noun-phrases in a sentence into references into the map. Thus, the map is effectively infinite by being implicitly self-extending upon demand.

So, before forging onward and adding new properties to the map and to TOUR, let us take stock of where we are. The representations for PLACEs and PATHs, which make up the map, have been described. The instructions to TOUR which make up route programs have been presented. TOUR and the "You Are Here" pointer have been described and their properties discussed. In addition, the interface to a problem-solver has been discussed. What has been presented thus far is sufficient to support an example of assimilation of knowledge into the map from a route description.

A mental map made from the representation as presented so far would consist primarily of stored routes, each quite well specified, with some topological knowledge about the PATHs and PLACEs involved, but few if any powerful problem-solving methods. Such a mental map would serve a person well in a situation that did not require knowledge of a huge geography, or rapid mastery or discovery of new routes.

TOUR IN OPERATION

Here I present an example of the operation of the TOUR machine on a simple route, to show how information is extracted from the route program and assimilated into the PLACEs and PATHs that make up the map. The example is very simple indeed, so it is important to notice that the assimilation methods used are independent of the size of the data base. They use only information directly accessible from the current instruction and the "You Are Here" pointer, and thus would work in exactly the same way within a much larger map.

Keep in mind, also, the distinction between TOUR, the model of spatial cognition, and MAPS, the program which is an example of it. TOUR is a mathematical abstraction which is used to describe states of partial knowledge about space. MAPS is a computer program (Note 5) which acts as a concrete model of TOUR. Some of the details described below, such as the representation of PLACEs and PATHs as property lists, are facts about MAPS, and could be done differently in some other model of TOUR.

For this example of assimilation, MAPS will begin with a very limited knowledge of part of Cambridge, Massachusetts, consisting of Harvard Square, Central Square, and MIT on Massachusetts Avenue (hereafter "Mass Ave"). It begins the example located at the intersection of Broadway and Prospect Street, whose relation with Mass Ave is unknown. This can be graphically shown as below:

```
                    B+PS


HS ---------------- CS ---------------- MIT
```

The initial state of the map for this example is listed below:

```
PATH1:                                  PLACE1:
        NAME:  Mass Ave                         NAME:  Harvard Square
        ROW:  (PLACE1 PLACE2 PLACE3)            ON:  (PATH1)
        EXTRAS:  nil                            STAR:  nil

PATH2:                                  PLACE2:
        NAME:  Broadway                         NAME:  Central Square
        ROW: (PLACE4)                           ON:  (PATH1)
        EXTRAS:  nil                            STAR:  ((0. PATH1 -1)
                                                        (180. PATH1 +1))
PATH3:
        NAME:  Prospect Street          PLACE3:
        ROW:  (PLACE4)                          NAME:  MIT
        EXTRAS:  nil                            ON:  (PATH1)
                                                STAR:  nil

                                        PLACE4:
                                                NAME:  nil
                                                ON:  (PATH2 PATH3)
                                                STAR:  nil
```

These descriptions are implemented as property lists with access functions which maintain the structure of the property values and can make limited deductions. A PATH description contains a partially ordered set of PLACEs under its ROW and EXTRAS properties. The largest totally ordered subset is under ROW, and additional fragments of the order are kept under EXTRAS. The access functions in this case attempt to merge the fragments into a total order whenever new information is provided.

The ON property of a PLACE description is a list of the PATHs which that PLACE is on. The STAR property represents the local geometry of the intersection. The value of the STAR property is a list of triples, each of which specifies a heading (in degrees) with respect to the local set of coordinates, followed by the PATH and DIRECTION having that heading. The DIRECTION on a PATH is specified as +1 or -1, and means "with" or "against" the order contained in that PATH's ROW property. Thus it is a conventional attribute of a PATH, and need not correspond to any global property. The example will begin with the "You Are Here"

pointer at the intersection of Broadway and Prospect Street, on Prospect Street, with an unknown DIRECTION.

```
YOU ARE HERE:
        PLACE:  PLACE4
        PATH:  PATH3
        DIRECTION:  nil
```

The street directions

This example will follow the knowledge states involved in the assimilation of the route description given below. A new piece of information in a description will be underlined. A recapitulation with minor variations will illustrate further properties. (Putnam Circle, which MAPS does not yet know, is in fact located on Mass Ave between Harvard Square and Central Square.)

```
"Take Prospect Street to Central Square."                              (1)
"Turn right."                                                          (2)
"Take Mass Ave to Putnam Circle."                                      (3)
```

MAPS receives street directions in a subset of English. They are translated in a straight-forward way to instructions for the TOUR machine. The two TOUR instructions that we shall be concerned with here are GO-TO and TURN. Thus, sentences (1) - (3) are translated into the internal instructions below. The amount of a turn is given as the resulting heading in an egocentric set of coordinates.

```
(GO-TO nil PLACE2 PATH3 nil)                                           (4)
(TURN nil nil nil 270. nil nil)                                        (5)
(GO-TO nil PLACE5 PATH1 nil)                                           (6)
```

while a new place-description is created for instruction (6):

```
PLACE5:
        NAME:  Putnam Circle
        ON:  nil
        STAR:  nil
```

Assimilation methods

Assimilation of the route-description by TOUR results in a three-way interaction among the instructions, the "You Are Here" pointer, and the map. The result is to move the "You Are Here" pointer along the path described in the instructions, adding new information to the map and to the route-description wherever possible. The elaborated route-description is saved, along with the additions to the map. The methods used in executing each instruction are given below, with the direction of transfer of knowledge, in each case.

A. Information from the "You Are Here" pointer can supply omitted initial conditions in the instruction. ["You Are Here" pointer ==> instruction]

B. Information in the instruction is sent to the relevant PATH and PLACE descriptions in the map. This information includes topological connections, ordering information on a PATH, and knowledge about the local geometry of a PLACE. [instruction ==> map]

C. The action specified by that instruction is performed, changing the "You Are Here" pointer. The map may deduce omitted information about the result of an action from what is already in the map: the street turned onto, the direction down the street, or the turn required. This inferred knowledge updates the route-description as well as the "You Are Here" pointer. [map ==> "You Are Here" pointer; map ==> instruction]

The first sentence

Sentence (1) is translated into instruction (4):

```
"Take Prospect Street to Central Square."                                    (1)
(GO-TO nil PLACE2 PATH3 nil)                                                 (4)
```

Method A provides information about the current position, producing:

```
(GO-TO PLACE4 PLACE2 PATH3 nil)                                              (4.1)
```

When executing a GO-TO instruction, Method B would like to send order information about PLACE2 and PLACE4 to PATH3. The instruction, however, does not specify their relative order, so the two fragments (PLACE2) and (PLACE4) are sent to PATH3. PATH3, via its access functions, can now define an order on the places it knows about. PLACE2 and PLACE4 are also told that PATH3 passes through them. The changed descriptions in the map are:

```
PATH3:                                    PLACE2:
        NAME:  Prospect Street                    NAME:  Central Square
        ROW:  (PLACE4 PLACE2)                     ON:  (PATH1 PATH3)
        EXTRAS:  nil                              STAR:  ((0. PATH1 -1)
                                                          (180. PATH1 +1))
```

PATH3 now has a defined direction, so Method C can deduce the direction of travel missing from instruction (4.1), to produce:

```
(GO-TO PLACE4 PLACE2 PATH3 +1)                                               (4.2)
```

and the "You Are Here" pointer becomes:

```
YOU ARE HERE:
        PLACE:  PLACE2
        PATH:  PATH3
        DIRECTION:  +1
```

(See note 6.)

The second sentence

The next sentence is:

```
"Turn right."                                                    (2)
(TURN nil nil nil 270. nil nil)                                  (5)
```

Method A supplies the initial state from the "You Are Here" pointer to produce:

```
(TURN PLACE2 PATH3 +1 270. nil nil)                              (5.1)
```

Method B adds no new information to the map: the connection between PLACE2 and PATH3 is already known, and there is insufficient information in (5.1) to add anything to the STAR property on PLACE2. When Method C performs the TURN instruction, however, it jumps to a conclusion about the street turned onto. The ON property of PLACE2 contains PATH1 and PATH3, and since the turn is made from PATH3, the destination of the turn is assumed to be PATH1. Not enough information is available to deduce the direction on PATH1. The result of this is:

```
(TURN PLACE2 PATH3 +1 270. PATH1 nil)                            (5.2)
```

and

```
YOU ARE HERE:
        PLACE:  PLACE2
        PATH:  PATH1
        DIRECTION:  nil
```

The third sentence

```
"Take Mass Ave to Putnam Circle."                                (3)
(GO-TO nil PLACE5 PATH1 nil)                                     (6)
```

Again, Method A provides the initial position, and confirms that we are on the right street. This produces:

```
(GO-TO PLACE2 PLACE5 PATH1 nil)                                  (6.1)
```

Exactly as in the first sentence, Method B can send only minimal order information to PATH1. This corresponds to knowing that Putnam Circle is on Mass Ave, but not knowing where. The result on the map is:

```
PATH1:                                      PLACE5:
        NAME:  Mass Ave                             NAME:  Putnam Circle
        ROW:  (PLACE1 PLACE2 PLACE3)                ON:  PATH1
        EXTRAS:  ((PLACE5))                         STAR:  nil
```

Finally, Method C moves the "You Are Here" pointer, but cannot deduce the direction of travel for this instruction.

```
YOU ARE HERE:
        PLACE:  PLACE5
        PATH:  PATH1
        DIRECTION:  nil
```

The route-description, extended by the inferences that have been made by TOUR, is saved for later use. It is indexed under the source and destination, and can be retrieved when that route is needed again. The final version of the route-description is:

```
(GO-TO PLACE4 PLACE2 PATH3 +1)                          (4.2)
(TURN PLACE2 PATH3 +1 270. PATH1 nil)                   (5.2)
(GO-TO PLACE2 PLACE5 PATH1 nil)                         (6.1)
```

Because of the local nature of the inferences, and the one-pass assimilation, more information can often be extracted on subsequent passes over the same route. For example, this route-description includes the fact that, wherever Putnam Circle is located on Mass Ave, a right turn at Central Square from Prospect Street points you in that direction. This fact is not captured by the partial order in PATH1. A later assimilation of this route may be able to extract that information for the map, once the local geometry of PLACE2 is more fully specified. This possibility is examined in the second alternative below.

Further Examples

It will illuminate the execution of a route by TOUR to discuss briefly a couple of alternate versions of the above example, and show how the result is different. First, consider the same route-description, but with sentence (3) replaced by:

```
"Take Mass Ave to Harvard Square."                                    (3.1)
```

Since the destination was previously known, Method C can deduce the direction of travel from PATH1.

```
(GO-TO PLACE4 PLACE2 PATH3 +1)                                        (4.2)
(TURN PLACE2 PATH3 +1 270. PATH1 nil)                                 (5.2)
(GO-TO PLACE2 PLACE1 PATH1 -1)                                        (6.2)
```

The next time this stored route-description is executed, the result of the TURN in (5.2) can be found by looking ahead to the following GO-TO.

```
(TURN PLACE2 PATH3 +1 270. PATH1 -1)                                  (5.3)
```

The TURN is now fully specified, so the next time Method B is applied, on a third pass over the route, the local geometry of PLACE2 can be augmented:

```
PLACE2:
        NAME:  Central Square
        ON:  (PATH1 PATH3)
        STAR:  ((0. PATH1 -1)
                (90. PATH3 -1)
                (180. PATH1 +1))
```

This additional knowledge about PLACE2 will increase the power of Method C to deduce the result of an incompletely specified TURN at PLACE2.

Consider now a second alternative to the original example. Assume that sentence (2) were:

```
"Turn right onto Mass Ave toward Harvard Square."                     (2.1)
```

Thus, after sentence (2.1),

```
YOU ARE HERE:
        PLACE:  PLACE2
        PATH:  PATH1
        DIRECTION:  -1
```

and

```
(TURN PLACE2 PATH3 +1 270. PATH1 -1)                                   (5.4)
```

and PLACE2 adds an item to its local geometry as described immediately above. Now, when instruction (6) is executed, Method A can fill it in completely:

```
(GO-TO PLACE2 PLACE5 PATH1 -1)                                         (6.3)
```

This, in turn, will allow Method B to send the ordered pair (PLACE5 PLACE2) to PATH1.

```
PATH1:
        NAME:  Mass Ave
        ROW:  (PLACE1 PLACE2 PLACE3)
        EXTRAS:  ((PLACE5 PLACE2))
```

This is a much more elaborately specified partial order, ripe for the addition of the item (PLACE1 PLACE5), which will allow the order to be linearized once more.

Summary of the example

This simple example has shown the way TOUR takes an underspecified route-description, and executes it to drive the "You Are Here" pointer through the map. The assimilation methods by which TOUR transfers information between several representations are shown. In addition, the example shows how small variations in the correspondence between the route-description and the current contents of the map make a large difference in the amount of knowledge that can be assimilated at that time. However, nothing from the route-description is discarded: it simply waits in the relatively inaccessible route representation until the map is ready to receive it.

Two additional important facts are worth emphasizing. The first is that the assimilation

process uses only locally available information, with no search of the map. Thus, assimilation proceeds independently of the size of the database containing the map. Second, assuming that the route description is correct, the description in the map may be partially specified, but it will not be wrong. The well-known inaccuracies in people's mental maps lie in the more global kinds of description which are presented below. In those cases, global conclusions are inferred from local observations, so they are quite vulnerable to error.

ORIENTATION

The most conspicuous lack in the description of spatial knowledge presented thus far is any notion of a "sense of direction." The only notion of direction that appears in the map thus far comes from the local geometry of a PLACE, which defines the heading of a (PATH DIRECTION) pair as it leaves that PLACE, with respect to an arbitrary, local coordinate system. Suppose, however, that a set of PLACEs shared the same (still arbitrary) coordinate system. Then headings from different PLACEs could be made comparable by being defined with respect to a common coordinate system. Add to this the notion of distance along a PATH, and we can define relative positions between pairs of PLACEs.

We can use this shared coordinate system to capture some of the knowledge involved in a "sense of direction." Let us define an ORIENTATION-FRAME as a set of PLACEs, along with a system of coordinates and a metric, with respect to which positions (distances and headings) can be defined between pairs of PLACEs. Just as partial knowledge of one-dimensional position is represented by the partial order in a PATH, partial knowledge of two-dimensional position can be represented by allowing an ORIENTATION-FRAME to know the relative positions of only some pairs of its PLACEs.

Another important aspect of partial knowledge of relative position is the idea of non-comparable positions. It is possible to have a number of different ORIENTATION-FRAMEs, within which the relative positions of PLACEs are defined. However, the relation between the different ORIENTATION-FRAMEs may be unknown, so that relative positions (especially headings) are not comparable between ORIENTATION-FRAMEs.

For example, consider two familiar parts of the Boston area: Harvard Square in Cambridge, and Government Center in Boston. It is very common for a person to have a good sense of direction within each of these areas without knowing the relation (relative position and orientation) between them. In this case, he has an ORIENTATION-FRAME for each area, which defines relative positions for many pairs of PLACEs in the area. However, the connection between the ORIENTATION-FRAMEs, if it exists at all, is very indirect. When asked about this relation, virtually everyone finds the problem difficult, and most give an answer which is at least partially wrong. (Note 7.)

The observation of distances and headings between pairs of PLACEs is simulated in the TOUR model by allowing a GO-TO on a straight PATH to obtain the distance between two PLACEs, and by providing a NOTICE instruction for TOUR, which contains the distance and heading of a "visible" landmark from a particular PLACE. The NOTICE instruction simulates visual observations: If the "You Are Here" pointer is at a given PLACE, PATH, and DIRECTION, then the remote PLACE can be observed at a certain heading and distance. The heading is defined with respect to an egocentric set of coordinates in which the heading of the observer is 0. In the MAPS program, a NOTICE instruction is the result of a verbal statement. The format of this instruction is:

```
(NOTICE <at-place> <on-path> <in-direction>
        <remote-place> <distance> <heading>).
```

For lack of an adequate theory of partial metrical knowledge, all distance and heading measurements are assumed to be accurate and precise. This assumption is demonstrably false for humans.

Having augmented the map with ORIENTATION-FRAMEs, the "You Are Here" pointer may

include, as part of its knowledge of the current position, a HEADING defined with respect to the current ORIENTATION-FRAME. This can allow us to compute the direction of other PLACEs in the same ORIENTATION-FRAME from a given position. It also means that there are two distinct kinds of orientation for the "You Are Here" pointer: DIRECTION, which is a one-dimensional orientation on the current PATH; and HEADING, which is a two-dimensional orientation with respect to the current ORIENTATION-FRAME. Like other parts of the "You Are Here" pointer, the current ORIENTATION-FRAME and HEADING may be unknown.

The relative positions of PLACEs which are not directly measurable can be computed, either using triangulation, or through a device called the IKON. The IKON performs a simple "visualization" of spatial relations, allowing a skeleton of known relations to determine the relative positions of many PLACEs. Its use is in some ways similar to the use of a piece of paper to draw a map. The paper acts as an external memory which allows certain global properties to be inferred, and global inconsistencies detected, which would not have been possible with the purely symbolic descriptions.

For example, if the relative positions of PLACEs A and B, and of B and C, are known by the same ORIENTATION-FRAME, then the relative position of A and C can be read off the IKON. Inconsistencies can be detected when two different PLACEs are drawn on top of each other, or when the relative position of two PLACEs according to the IKON is different from that in the ORIENTATION-FRAME. In both cases, this indicates that the geographic knowledge, while locally consistent and apparently correct, is globally wrong. [This kind of global inconsistency will result in wrong answers from some of the problem solving procedures. Most of the processing will be unaffected, however, and an inconsistency which would be troublesome to remove may simply be left in.] Facilities for analyzing these inconsistencies to isolate and

correct errors in geographical knowledge have not been implemented, so the use of the IKON is still poorly understood.

Relative position information can be used in a number of ways. It can, of course, be used to answer direct questions. It is also useful in problem-solving, where knowledge of the position of the goal can guide exploration even when a complete route cannot be found. A most important use of position information is in orienting oneself. The ORIENTATION-FRAME is an abstract system of coordinates, so one must continually be computing where it is with respect to the observable environment. A NOTICE instruction at a particular point on a route allows TOUR to determine its HEADING with respect to that ORIENTATION-FRAME from the observed (egocentric) direction of a distant PLACE.

For example, in Boston many people orient themselves with respect to the Prudential Building and the John Hancock Building. These landmarks are visible from a great distance, and allow a person to deduce his own heading and position with respect to the ORIENTATION-FRAME which represents his position kowledge. Thus he can estimate the egocentric direction of PLACEs which are not visible.

The ORIENTATION-FRAME is an arbitrary frame of reference for defining relative positions of PLACEs. In human development, a later stage brings the recognition that there is a distinguished conventional ORIENTATION-FRAME: the cardinal directions. This globally-available ORIENTATION-FRAME allows useful relations to be stated between directions in remote areas that would otherwise remain unrelated. The developmental sequence, from topological to local orientation to conventional orientation, has been documented in the psychological literature, beginning with Piaget and Inhelder (1967).

ABSTRACTION

The model of the map that we have developed so far consists of a network-like collection of PLACEs and PATHs, and a number of ORIENTATION-FRAMEs which allow two-dimensional position to be defined among sets of PLACEs. There is, however, another important feature of human spatial knowledge that is not described by this kind of map. Abstraction, in spatial cognition, is the ability to use a schematized version of the map, omitting much of the known detail, in order to solve problems efficiently. We can modify the current definition of the map to encompass this phenomenon by allowing the map to be made up of several disconnected components, each describing the same geography at different levels of detail. This requires us also to specify how a correspondence is represented between different components of the map.

It is likely that as the geography of an area is learned, the map consists of disconnected parts, representing different regions that are known, but not connected. To provide an abstract description, two components must describe the same geographical area at different levels of detail. It thus seems reasonable that the ability to represent abstraction develops from the ability to represent unrelated areas. However, it remains an open question how the additional relations are learned.

Two components of the map describing the same geography will have different levels of detail, so I will refer to the more schematic component as the "higher" one, and the more detailed component as the "lower" one. A higher PLACE (i.e. a PLACE in a higher component) corresponds to a set of lower PLACEs, called a REGION. A higher PATH corresponds to a

lower PATH (or several of them), and provides an order on a set of higher PLACEs.

The use of the higher components of the map is to simplify problem-solving. The problem of finding a route between two PLACEs in the lower component may be very complex, but correspond to a much simpler problem in a higher component. This leaves the difficulty that the proposed route begins and ends in the wrong component, so TOUR needs special instructions to move the "You Are Here" pointer between corresponding positions in different components.

However, a higher PLACE corresponds to many lower PLACEs. How can positions correspond between two components when they differ so in amount of detail? The answer is that the higher components are built so that a fully specified "You Are Here" pointer (PLACE PATH DIRECTION) in the higher component can be mapped onto a single fully specified "You Are Here" pointer in the lower component. This correspondence is implemented as a property of the higher PATH, indexed under the PLACE and the DIRECTION. There is an analogy between this correspondence and that from the exits of a limited-access highway to the surrounding surface roads.

For example, consider the Mass Pike as a higher PATH, containing as PLACEs Boston, Cambridge, Newton, and Framingham. In this higher component, the "You Are Here" pointer has as its state:

```
(Cambridge, Mass Pike, eastbound)
```

which corresponds to a "lower" value of the "You Are Here" pointer at a PLACE near the tollbooth, where a choice must be made between a PATH leading to River Street in Cambridge, and one leading to Cambridge Street in Allston.

Essentially without physical motion, TOUR can change its "point of view" or "level of abstraction" by changing the component where the "You Are Here" pointer is, if the appropriate mapping is known. The UP and DOWN instructions specify corresponding states of the "You Are Here" pointer in separate components, and instruct TOUR to move from one to the other. UP specifies motion from a lower component to a higher one, and DOWN specifies the reverse.

```
(UP <lower-place> <lower-path> <lower-direction>
        <higher-place> <higher-path> <higher-direction>)
(DOWN <higher-place> <higher-path> <higher-direction>
        <lower-place> <lower-path> <lower-direction>).
```

A curious feature of these abstractions is that, although a TURN in a higher component leaves the "You Are Here" pointer at the same higher PLACE, it can correspond to travelling a significant distance in the lower component. For example, consider the higher route:

```
Take I-93 to Boston.
Then take the Mass Pike to Cambridge.
```

TOUR will fill in a TURN instruction at the PLACE Boston, changing the "You Are Here" pointer from I-93 to the Mass Pike. This requires a route to be found in the lower component for accomplishing the TURN.

A question that remains open is how these abstract representations of the geography are created automatically. Once the higher components exist and routes within them can be found and followed, they can undergo the same kinds of gradual refinements as the lower components. However, it is not clear how to create these descriptions from local experience.

GLOBAL DESCRIPTIONS

At this point, the map is a network-like collection of PLACEs and PATHs, a set of ORIENTATION-FRAMEs with positions defined between pairs of PLACEs, and separate components so that the same descriptive techniques can be applied to the same geography at several levels of abstraction. In addition, the TOUR machine can allow routes to be specified in various incomplete ways, filling in the gaps by posing them as problems to be solved, or by extracting the necessary information from the map.

There are, however, additional ways of describing the space to capture more global (or "gestalt") properties. One category of such properties includes knowledge about streets being parallel, and regions being laid out as rectangular grids. The rectangular grid lends itself to powerful route-finding techniques, so it is a very useful way to describe a region. This makes it particularly popular, and many people apply it well beyond its legitimate domain of applicability, leading to the most frequent errors in people's mental maps. Another category of such global properties includes the outline shapes of regions. The outline of a region can act as a boundary, with a small number of "gates" allowing travel across. This simplifies problem-solving by specifying a limited set of potential intermediate points for a route that crosses a region boundary.

A REGION is a set of PLACEs and PATHs to which one or more of these global descriptions is applied. Thus, route-finding within a REGION is often simplified by the availability of the powerful problem-solving methods permitted by the global descriptions. At the same time, REGIONs lend themselves to abstraction by correspondence with a single higher PLACE. This

simplifies route-finding between REGIONs as well.

A REGION which is described as a rectangular grid contains two bundles of parallel PATHs, each bundle partially ordered. An ORIENTATION-FRAME is associated with such a REGION to define the headings of the bundles, and the heading of the order of the PATHs within each bundle. A route-finding problem within a rectangular REGION is first expressed as the problem of getting from one grid intersection to another. Then it can be reduced to two independent problems to be solved in parallel, one in each bundle of PATHs. The goal of each problem is to reach a particular PATH. When both are achieved, the desired intersection has been reached. The power of the rectangular grid description is the assurance that these subgoals can be pursued independently to solve the original problem.

The description of a REGION as a rectangular grid can be built up from locally observable pieces of information. When TOUR executes a GO-TO instruction, it sends the source and destination PLACEs and the connecting PATH to a separate process which attempts to build a rectangular grid description. If a number of conditions are met by that piece of the map, it is added to a growing REGION and incorporated into the rectangular grid. A REGION will grow by this local means until it is bounded by regions that fail to meet the conditions. Often such a boundary is formed by another rectangular grid at a different orientation. Neither grid can propagate across their common boundary. An example of this is Market Street in San Francisco, which separates two grids at a 45 degree angle to each other, and which is a source of geographical confusion to both residents and tourists.

The boundary description of a REGION does not seem so amenable to local propagation, unfortunately. This description apparently goes through two stages in its development. The

first identifies a collection of geographical features as "edges" which form a boundary to the REGION, with "gates" for travel across. The boundary acts purely as a container for the REGION, without a shape being represented. The second stage describes the shape of the boundary, and seems to call on visual knowledge for shape descriptions. It may also be that an accurate shape description depends on some sort of global visual observation, either from a printed map or from an altitude.

The Charles River acts as a boundary for both Boston and Cambridge. The few bridges which cross it provide a small set of intermediate points on routes being sought between the two cities, even when the shape of the river is not known. A correct description of the river's shape can be useful in route-finding, but the mistaken assumption that the Charles is straight is a popular source of geographical confusion in the Boston area.

The rectangular grid description, with problem-solving and region-growing mechanisms, has been implemented in the MAPS program. The boundary description has not been implemented.

PROBLEM SOLVING

A number of problem-solving mechanisms have been referred to in the description of TOUR thus far. They play an important role in the ability of TOUR to accept underspecified routes by calling on the problem-solver to fill in the gaps. In fact, verbal directions given by people seem to be designed to provide, not complete routes, but the key facts that leave only easily solvable problems. Thus, it is important not only to have a problem-solver, but to have an accurate description of its capabilities.

The simplest and most straightforward problem-solving technique is to use previously known routes. These are indexed by source and destination in the map, and are chosen to cover the most frequent needs. When a new one is needed, rather than solving the problem by examining the map, the person can ask for a verbal route-description from a reliable informant, assimilate it with the TOUR machine as shown in the example, then store it to be used when needed, and filled in and clarified by TOUR on successive trips. If the demands on one's geographical knowledge are light, this is a very practical strategy.

The next kind of useful problem-solving is the simple network search on the map. When the set of possibilities can be pruned to a very small number, this can often be effective. There are several methods of pruning that use different aspects of the map. The first is to use information in the current ORIENTATION-FRAME to suggest intermediate points located between the source and the goal. The second is to rephrase the problem as one in a more abstract component of the map, with a smaller set of PLACEs and PATHs, spread over a larger area.

The most sophisticated problem-solving techniques use global properties of the map as described in some REGION. These were discussed in the last section along with the presentation of those region descriptions.

In general, it seems that people who are good geographical problem-solvers use all of these methods in the appropriate circumstances, but that the first method, that of knowing the route already, is the most effective. Ability to navigate in a geographical area seems to be primarily dependent on the amount of knowledge represented about an area, and only secondarily on the techniques used to solve problems. On the other hand, the organization that imparts the most problem-solving power is the rectangular grid, even when it is wrong in minor ways. This doubtless accounts for its over-enthusiastic application. Most of the obvious inaccuracies in people's mental maps seem to derive from the incorrect assumption of straight streets and rectangular grid patterns. {Note 7.}

PARAMETERS OF VARIATION

One of the most striking features of human spatial cognition is the range of individual variation, both in the amount and the kind of knowledge represented. Although it is hardly surprising that any reasonably complex computational model could support enormous amounts of variation, it is interesting to examine exactly where in the model it takes place, and how it can vary.

A simple production system control structure (Newell and Simon 1972) is used in several places in the TOUR model. A problem to be solved or an action to be taken is described. There are a number of simple procedures which can recognize and act on special cases of the situation described. They are activated in sequence, and the first one which recognizes its special case handles the situation. Variation can take place in such a system, not by the value of a numerical parameter, but by the selection and ordering of the sequence of procedures to be applied to a particular problem.

These simple production systems occur most importantly in three places in the TOUR model. First, the TOUR machine itself operates by means of a production system which examines each instruction along with the "You Are Here" pointer and the relevant portions of the map. Second, a problem is solved by a production system of potential methods for solution. Third, and most important, the global descriptions of the geography are constructed by a production system of specialists who examine portions of the map suggested by TOUR instructions with special properties.

The largest part of the qualitative variation among people in their styles of spatial cognition seems to come from the presence or absence of these global descriptions. Among those who do represent global descriptions, there is a very noticeable variation according to how carefully they check for the legitimate applicability of that description. I have interviewed a subject who was even willing to sacrifice topological connections to preserve the description of (what he believed to be) a rectangular grid. {Note 7.}

The presence or absence of global descriptions of the geography certainly affects the selection of methods available for problem solving, since many of those are dependent on particular global descriptions. A great deal of empirical study must be done to clarify the exact nature of variation in human spatial cognition, and how that relates to the qualitative parameters of variation in the TOUR model.

THE PSYCHOLOGICAL LITERATURE

There has been a considerable amount of interest among psychologists in spatial cognition of various kinds. Downs and Stea (1973) and Siegel and White (1975) both include extensive bibliographies. Most of this section is devoted to a detailed comparison of the TOUR model with the summary of the psychological evidence done by Siegel and White (1975).

After a number of interesting but isolated pieces of research in the first half of the century, recent interest in spatial cognition was sparked by "The Image of the City" (1960) in which Kevin Lynch, coming from an urban planning tradition, explored the phenomena of subjective perception of the city. This delightful book inspired a flood of further research by urban planners, geographers, and cognitive psychologists.

This surge of interest in spatial cognition resulted in a profusion of different viewpoints on and definitions of the research area. Downs and Stea (1973) collect a number of papers showing the diversity and vitality of the emerging field. A review article by Hart and Moore (1973) surveys the work of psychologists looking at the development of spatial cognition in children. Most of the current work is based on the theories and methodologies of Piaget and Inhelder (1967). An excellent example of modern developmental theory following Piaget and Inhelder is the work of Moore (1972).

These developmental studies, however, looked at many different kinds of spatial cognition. Siegel and White (1975) focus specifically on work dealing with spatial representations of large-scale environments. The model of spatial knowledge that Siegel and White synthesize

from their survey is very similar to the model that I had developed, and was responsible for a number of important refinements leading to the TOUR model.

In the paragraphs below, I quote extensively from Siegel and White, to summarize their theory of spatial representations of large-scale environments. I also discuss the differences between their theory and the TOUR model. The TOUR model is intended to provide a precise framework for stating a theory like theirs, suggesting useful modifications, and proposing possible experiments.

> *"The terminology has a tendency to suggest pictures or maps, but a variety of research indicates that 'images' are not 'maps' and are often not even map-like.*
>
> *1. The representations are typically fragmented. Areas of considerable detail are linked with areas having little or no detail and are often separate from one another. ...*
>
> *2. The representations are often distorted. ...even topological and projective relations are often not retained.*
>
> *3. The representations are often several separate, but interlaced representations of smaller chunks of the environment. ...*
>
> *4. The representations do not need to be entirely visual. ..." (p. 21)*
>
> *"Most theorists essentially agree that landmarks and routes are the predominant elements of spatial representations. From a logical point of view it can be argued that landmarks and routes are perhaps the necessary and sufficient elements for 'minimal' representations that allow 'way-finding' to occur.*
>
> *... Landmarks are unique patterns of perceptual events at a specific location, they are predominantly visual for human adults, they are the strategic foci to and from which one travels, and they are used as proximate or intermediate course-maintaining devices.*
>
> *... Routes are nonstereotypic sensorimotor routines for which one has expectations about landmarks and other decision points. ... They represent habitual lines of movement and familiar lines of travel, and thus they constitute a first-order, 'enactive' representation of the terrain. One can conceive, then, of the environment consisting of potential landmarks connected by potential routes. One can picture a spatial representation as landmarks (visual 'pegs') connected by routes (sensorimotor 'lines'), to some extent guided by sequence learning.*

... In addition to landmarks and routes, a third useful and often-present element of a spatial representation is constituted by gestalt knowledge. Knowledge of configuration gives something more than a minimal map. It is a sophisticated wrinkle that gives its owner an advantage in way-finding and organizing experience. There seem to be at least three types of such knowledge of configurations: A perceived outline of a terrain (e.g., the outline of a United States map); a graphic skeleton (e.g., a schematic portrayal, a spatial representation of London as a set of routes leading from a diagrammatic image of the subway system); and a figurative metaphor (e.g., the 'boot' of Italy). These 'configurations' enhance way-finding, and they may be a necessary condition for inventions of new routes. We would argue that all spatial representations are functionally "landmarks-connected-by-routes," but that there are varying degrees of integration or gestaltness of the spatial representation." (pp. 23-24)

A number of points of comparison are in order here. A PLACE in the TOUR model represents only part of the knowledge that Siegel and White associate with a landmark, omitting the complex of immediate sensory information. This sensory image allows a human to identify the landmark he is at without knowing how he got there. Although this knowledge is important to humans, I believe that it can be factored out of the TOUR model and treated as knowledge of a different kind for purposes of describing spatial cognition. This factoring is based on the assumption that the sensory image of a PLACE is important only in identifying the current PLACE, and does not play a major role in problem solving or assimilation. A more refined model will have to reexamine this assumption.

Siegel and White's routes are clearly very similar to TOUR programs. However, they omit PATHs from their representation of the geography. It has seemed necessary to me to separate the function of representing the order of PLACEs on a PATH (as a one-dimensional geographical feature) from the function of representing the procedural description of a route used to travel from one place to another. This is basic to the distinction between the map, which is a description of the geography, and the routes, which describe activities within that geography.

The TOUR model uses the ORIENTATION-FRAME to capture the course-maintaining function of landmarks. In addition, the ORIENTATION-FRAME represents partial knowledge of relative position, an important part of human spatial cognition even aside from the course-maintaining function. Only in passing do Siegel and White recognize the possibility of multiple ORIENTATION-FRAMEs, or "minimaps," allowing potentially incomparable relative positions to be defined.

The graphic skeleton, which is one of the gestalt configurations proposed by Siegel and White, corresponds quite well to the abstract, schematic component of the map. The other gestalt configurations strongly resemble the global descriptions which I have sketched briefly.

> *"One can consider a route as a conventional sensorimotor system. Although we are not in the position of providing a formal analysis of what this system is, it is possible to point to likely elements or characteristics of what the system must have.*
>
> *1. A route must involve a sequence of decisions--generally, changes in heading. In several attempts to set up models of route making, we have been unable to project a reasonable kind of route learning without entering into the knowledge system some such entry as 'bearing' or 'heading.' Unless the organism steers its way through a route by a sequence of line-of-sight approach movements, a kind of process that seems unlikely for all route learning, the organism must compute its decisions in terms of an orientation of the direction of the organism with regard to some feature of the environment.*
>
> *2. The knowledge of a route could then conceivably exist through a kind of serial learning, a memorized series of decisions. However, it seems much more likely that a memorized route would be somewhat more akin to paired-associate learning, changes in bearing asociated with the arrival at 'stimulus' landmarks.*
>
> *3. Learning between landmarks is, to some extent, incidental and irrelevant except to the extent that intermediary landmarks serve as course-maintaining devices (and thus as landmarks associated with no change in heading). A conservative route learning system would then be, in effect, 'empty' between landmarks. All of the learning would be organized around the nodes of the decision system, the landmarks. In the adult's construction of a spatial representation, routes become scaled by landmarks in an ordinal and roughly interval sense." (pp. 28-29)*

In the TOUR model, there are two senses of "orientation" that are used: DIRECTION with respect to the one-dimensional order on a PATH, and HEADING with respect to the two-dimensional positions in an ORIENTATION-FRAME. DIRECTION orientation is defined with respect to local features of the map network, and does not imply knowledge of a global heading. The assimilation example shows how this kind of one-dimensional orientation can be acquired from local evidence.

> *"Once one locates oneself along a number of routes by a system of landmarks, these routes with termini become interrelated into a networklike assembly as a function of repeated experience, temporal integration, and sustained meaningfulness. Taxi drivers, for example, organize such networklike knowledge of American cities. They associate conditions of time, traffic, road conditions, etc., to components of the network. With such knowledge a taxi driver can plan a route through the maze of city streets which will be both as short as possible, and which will optimize speed of travel under prevailing conditions. ... Once routes with termini become interrelated into a networklike assembly, the gaps are gradually filled in [and become Lynch's (1960) 'districts,' 'edges,' and 'nodes']. The landmark-connected-by-routes spatial representation becomes more gestalt-like." (p. 30)*

> *"Landmarks and routes are the minimal elements of spatial representations; they lead to what has been termed route-representations. Survey-representations incorporate configurational elements (outlines, graphic skeletons, figurative metaphors) and may be the final derivative of very dense and richly interconnected and hierarchically organized route maps." (p. 45)*

The TOUR model specifies how knowledge from the route is incorporated into the map, into descriptions of PLACE, PATH, ORIENTATION-FRAME, and REGION, and then how the map is further described from a global point of view to assist in problem-solving. Siegel and White also observe this process, but their theory says only that some sort of union of the set of routes will constitute the global knowledge of the geography.

> *"With regard to children, the following sequence in the development of spatial representations of the large-scale environment has been identified in the research literature. The sequence is the sequence with which we have become familiar in discussions of adult learning. ...*

1. Landmarks are noticed and remembered. ...

2. Once landmarks are established, the child's acts are registered and accessed with reference to them. ...

3. Given landmarks, action-sequences, and organized decision systems about where to look next, the child forms clusters of landmarks and 'minimaps.' ...

4. A key issue in the development of spatial representation is the child's formation of some kind of objective frames of reference, and a concomitant organization of outside features into systems in space. Part of what is involved in this process appears to be a progressive differentiation of self-orientation from outside-orientation, and the development of a notion of objective bearing. ...

5. Survey maps appear as coordinations of routes within an objective frame of reference. That is, survey maps become possible only after both routes and an objective frame of reference exist." (pp. 37-45)

"... the development of children's spatial representations conforms to the 'Main Sequence' identified in adult learning: The process of going from landmarks, to route-maps, to survey-maps is a process of going from association to structure, and of deriving simultaneity from successivity." (p. 46)

This learning sequence is also apparent in the way that TOUR assimilates new information. This model, unfortunately, sheds little light on how the TOUR machine itself is learned during the development of the child. However, once its operation is precisely described, it may be possible to represent it in a learnable way.

TAKING STOCK

As is by now quite clear, this paper reports work in progress. This research will be more fully reported in Kuipers (1976). Most of the TOUR model has been implemented as the MAPS program, with some exceptions which will be discussed below. Although it is possible to describe the TOUR model as a mathematical abstraction, independent of any computer implementation, it would not have been possible (for me) to discover it. When attempting to create a model as complex and precisely specified as the TOUR model, it seems impossible to evaluate the consequences of technical design decisions without a computer implementation. The MAPS program has also allowed me to experiment with the model at various stages, to discover which aspects are important. I have clearly specified where the model does not have an implemented version, and where further development is required in the theory.

The first part of the problem we set ourselves was to identify and describe a class of human behavior to which a model could be addressed. This behavior is the assimilation of fragmentary pieces of information into a variety of representations which can support the kinds of problem-solving people do. [As we see, it is difficult to describe the phenomena we are interested in apart from the vocabulary of the model we will use to explain it.] The assimilation process consists of 1) maintaining a correspondence between the input route description and the map, 2) filling in missing information in the route description from what is known by the map, 3) filling in the map from new information provided in the route description, and 4) building higher-level representations by describing selected pieces of the map. The kinds of problems that people are thereby enabled to solve include finding new routes from one place to another, and orienting themselves with respect to the positions of

remote places. The greatest skill that people have in this domain is the assimilation of new information into appropriate representations, rather than deep and complicated problem-solving.

The second part of our original problem was to propose a model---a computational process---which explains these phenomena. I have described a representation for the map, consisting of PLACEs, PATHs, and ORIENTATION-FRAMEs, with separate components allowing several abstract descriptions of the same geography. In addition, REGIONs permit the application of global descriptions to subsets of the map. A problem solving process examines the map to propose routes between one PLACE and another, and another process attempts to build global descriptions and the REGIONs they apply to. The TOUR machine plays a central role in the interaction between parts of the model. In particular, it assimilates information from the route-description into the map, it uses information from the map to fill in omitted parts of instructions, and it calls the problem solver to provide missing parts of the route.

Thus, although some parts of the model are still somewhat incomplete, we have proposed a solution which appears to satisfy the conditions of our original problem. As mentioned above, there are several parts of the model which have not been fully implemented, and are therefore not as well understood as others.

A. Several kinds of problem-solving are implemented, but they lack a coherent framework of the kind that clarifies the rest of the model.

B. The methods for creating higher level descriptions are also poorly understood, although some have been implemented.

C. The IKON has not been interfaced to the rest of the TOUR model.

D. Metrical information is assumed to be accurate and precise, an assumption about people which is known to be false in interesting ways.

E. The way people use the sensory image of the geography deserves further study.

There are some aspects of people's mental maps of cities which have been omitted from the TOUR model in order to focus on parts of the model that apply more generally to cognition of large-scale spaces. Some of these will be added to MAPS eventually to make it a more interesting and useful computer program.

A. A street has two sides, and the relation "across the street from" is a useful one.

B. Streets and intersections have traffic limitations which should be taken into account: traffic lights, illegal turns, one-way streets, etc.

C. Proposed routes can often be evaluated for cost, usually by distance or time, but also for convenience, quality of road surface, character of neighborhood, and time of day.

The immediate experimental test to which the TOUR model can be subjected is to test the predictions about what kinds of partial knowledge are actually represented. For deeper results, it would be nice to examine the behavior of an individual in comprehension, problem-solving, or explanation. Unfortunately, there are serious methodological problems which must be solved before the TOUR model can be fitted to an individual so such an experiment can be set up. This custom-fitted model is required because overt behavior can vary so widely depending on small variations in the knowledge the subject has, how it is represented, and the particular methods he has for processing it. Each of these problems may itself require much experimental study.

A. Identify which parameters of variation in the model are parameters of individual variation, and which remain fixed across a culture, or across all humans.

B. Determine the settings of these parameters for the individual subject.

C. Determine exactly what knowledge the subject has about some geographical area, either by testing him, or by teaching him a carefully designed body of knowledge about an unfamiliar area.

This is certainly an awesome problem, but people are awesome in their complexity, with individual variation being a particularly striking feature. One contribution of the TOUR model is to provide a descriptive framework for stating the answers to these three problems.

Without that, we cannot even express the questions discussed below.

There are, of course, many aspects of this enormous and complex domain of knowledge which are not directly addressed by the TOUR model. I believe, however, that the TOUR model provides a framework which will allow these phenomena to be investigated with a clarity that was not previously possible.

A. There are interesting linguistic conventions for describing routes and relative positions.

B. The process of understanding verbal directions given by a fallible or dishonest speaker is different from the process of understanding "observational" (i.e. presumably correct) route directions.

C. People are able to detect inconsistencies in their knowledge, and replace them with more accurate information, under some circumstances.

D. One dramatic subjective event is the "flash of comprehension" where a person recognizes a single new geographical fact which reorganizes a large amount of his knowledge of an area.

E. People occasionally have separate descriptions of PLACEs which turn out later to describe the same physical place. This realization can trigger the "flash" mentioned above.

F. Spatial metaphors play a prominent role in memory and problem-solving in non-spatial domains.

An important assumption of this research is that the *representation* of knowledge, and the assimilation of new information into that representation, is fundamental to spatial cognition. The other aspects mentioned above are built on this fundamental basis, in the sense that they cannot be fruitfully investigated without a prior theory of representation. I have attempted to present the TOUR model as a precise way of describing the states of knowledge involved in large-scale spatial cognition. I hope that it will permit a kind of research that was not previously possible.

NOTES

{1. TOUR "machine"}

The TOUR machine is an abstract machine in the same sense that a Turing machine is. A Turing machine is a finite-state machine operating on an infinite tape. Typically the finite-state machine and its instruction set are simple, while most of the machine's complexity is represented by the contents of the tape. Similarly, TOUR is a simple processor with a small instruction set, operating on a potentially infinite and very complex map.

{2. LOGO analogy}

A useful analogy is to the LOGO programming language, which was initially developed to teach programming concepts to children. LOGO instructions drive a simple robot "turtle" around the floor. The state of the turtle is completely described by its position and its heading. The two basic instructions are:

```
FORWARD <number of units>
RIGHT <number of degrees>
```

The turtle also has a pen which can be raised and lowered to draw lines on the floor as it moves. These instructions are enough to allow children to draw many kinds of interesting pictures. Adding the ability to define procedures to be used as parts of other drawing programs makes LOGO a powerful tool for a child to use for creating pictures. The process of writing and debugging such drawing programs is believed to encourage the development of powerful ways of "thinking about thinking."

{3. Definitions}

A PLACE is a description of a zero-dimensional geographical object. It includes the PATHs which that PLACE is on, and a local geometry defining the radial headings of (PATH DIRECTION) pairs leaving the PLACE, with respect to an arbitrary coordinate system. The description need not be complete.

A PATH is a description of a one-dimensional geographical object. It includes a partial order on the set of PLACEs on that PATH.

The DIRECTION on a PATH can be +1 or -1, meaning "with" or "against" the order on the PATH.

An ORIENTATION-FRAME is a set of PLACEs, along with a system of coordinates with respect to which distances and directions can be defined between pairs of those PLACEs.

A HEADING is a direction (in degrees) defined with respect to an ORIENTATION-FRAME.

A REGION is a set of PLACEs, to which a global description can be applied; for example, nature of grid or shape of outline.

The map is a set of PLACEs, PATHs, ORIENTATION-FRAMEs, and REGIONs.

The "You Are Here" pointer describes the current position of the TOUR machine in the map. It includes the current PLACE, PATH, DIRECTION, ORIENTATION-FRAME, and HEADING. Some of these may be left unspecified.

{4. TOUR Instructions}

GO-TO instructs TOUR to move the "You Are Here" pointer from one PLACE to another on the same PATH, and specifies the direction of travel with respect to the PATH order.

```
(GO-TO <from-place> <to-place> <on-path> <in-direction>)
```

TURN instructs TOUR to move the "You Are Here" pointer from one (PATH DIRECTION) pair to another at the same PLACE, and specifies the amount of the turn.

```
(TURN <at-place> <from-path> <from-direction>
      <turn-amount> <to-path> <to-direction>)
```

The TAKE instruction refers to a route, and lets it be used as a sub-program for a longer route. This can give a route program a hierarchical structure, and it allows frequently-used routes to be shared.

```
(TAKE <route> <from-place> <to-place>)
```

The GET-TO instruction formulates a problem to be solved by the problem-solving component by specifying the state of the "You Are Here" pointer at source and destination. In case no

solution can be found, the problem-statement is left as an instruction in the route program, and TOUR continues from its destination.

```
(GET-TO <from-place> <from-path> <from-direction>
        <to-place> <to-path> <to-direction>)
```

The NOTICE instruction simulates visual observations: if the "You Are Here" pointer is at a given PLACE, PATH, and DIRECTION, then the remote PLACE can be observed at a certain heading and distance. The heading is defined with respect to an egocentric set of coordinates in which the heading of the observer is 0.

```
(NOTICE <at-place> <on-path> <in-direction>
        <remote-place> <distance> <heading>)
```

The UP and DOWN instructions specify corresponding states of the "You Are Here" pointer in separate components, and instruct TOUR to move from one to the other. UP specifies motion from a lower component to a higher one, and DOWN specifies the reverse.

```
(UP <lower-place> <lower-path> <lower-direction>
        <higher-place> <higher-path> <higher-direction>)
(DOWN <higher-place> <higher-path> <higher-direction>
        <lower-place> <lower-path> <lower-direction>)
```

{5. The MAPS program}

The MAPS program is written in MACLISP, and runs on the MIT AI Laboratory's PDP-10. It uses about 100K of memory, including the LISP interpreter. Its natural language input is provided by a small context-free grammar with 86 rules and a vocabulary of 123 words plus place-names which can be learned on the fly. It typically reads, parses, and processes a sentence in less than a second.

{6. Diagrams}

It is clear that diagrams would help an example like this. Unfortunately, existing graphical conventions for drawing maps make it difficult to represent partial geographical knowledge without committing ourselves to something actually false.

{7. Experimental results}

This experimental result is a qualitative conclusion drawn from my own interviews, which were not systematically designed or analyzed. Given the theoretical framework of the TOUR model, it would be valuable to do a more careful interview study.

References

Roger M. Downs and David Stea (eds.), "Image and Environment." Chicago: Aldine Publishing Company, 1973.

Thomas Gladwin, "East is a Big Bird." Cambridge: Harvard University Press, 1970.

Roger A. Hart and Gary T. Moore, The Development of Spatial Cognition: A Review, in [Downs & Stea 1973].

Benjamin Kuipers, The Representation of Knowledge of Large-Scale Space. Doctoral thesis, MIT Mathematics Department, forthcoming 1976.

Kevin Lynch, "The Image of the City." Cambridge: MIT Press, 1960.

Gary T. Moore, Elements of a genetic-structural theory of the development of environmental cognition, in W. J. Mitchell (Ed.), "Environmental Design: Research and Practice," v. 2, Los Angeles: University of California, 1972.

Allen Newell and Herbert A. Simon, "Human Problem Solving." Englewood Cliffs, NJ: Prentice-Hall, 1972.

Jean Piaget and Baerbel Inhelder, "The Child's Conception of Space." New York: Norton, 1967. (first published in French, 1948)

A. W. Siegel and S. H. White, The development of spatial representations of large-scale environments, in H. W. Reese (ed.), "Advances in Child Development and Behavior," v. 10, New York: Academic Press, 1975.